BLOG @ CACM

The *Communications* Web site, <http://cacm.acm.org>, features more than a dozen bloggers in the BLOG@CACM community. In each issue of *Communications*, we'll publish selected posts or excerpts.

Follow us on Twitter at <http://twitter.com/blogCACM>

DOI:10.1145/3276740

<http://cacm.acm.org/blogs/blog-cacm>

The Gap in CS, Mulling Irrational Exuberance

Carl Hewitt suggests computer science needs a reference resource, while Vijay Kumar decries intellectual dishonesty in technology forecasting.

Carl Hewitt
Computer Science
Encyclopedia
Can Fill a Gap

[https://cacm.acm.org/blogs/blog-](https://cacm.acm.org/blogs/blog-cacm/230860-computer-science-encyclopedia-can-fill-a-gap/fulltext)

[cacm/230860-computer-science-encyclopedia-can-fill-a-gap/fulltext](https://cacm.acm.org/blogs/blog-cacm/230860-computer-science-encyclopedia-can-fill-a-gap/fulltext)
September 5, 2018

There is an important gap in computer science (CS) education and professional collaboration that can be filled by a nonprofit online reputable, referenceable encyclopedia supported by appropriate professionally relevant advertising. The encyclopedia should be managed by a prestigious editorial board that appoints a hierarchy of editors to moderate articles. The editorial board can guarantee editorial independence from advertisers analogous to current professional practices for journals and conferences. Anyone would be allowed to register to submit suggestions and drafts to the editors. Access to articles would be free and available to all. The encyclopedia must establish procedures to be fair and inclusive on the basis of race, sex, religion, age,

disability, and national origin integrating content suitable for all from preschool to advanced researchers.

The encyclopedia could support interactive articles with videos, animations, and dynamic narrations. Within a decade, interactive content could be a requirement for most articles. Over time, the encyclopedia should be organized using ontological services supporting programmatic interfaces for a knowledge graph.

The encyclopedia should become a standard reference, a trustworthy professionally accountable educational resource for all. Currently, there is no online encyclopedia that can serve as the source of valid scientific references.

Our profession has the credibility and resources to create an encyclopedia to serve as the professional standard. Serving as a member of its editorial board could be a prestigious office for senior professionals to provide experience and judgment. Professional reputations could be enhanced by contributing to the encyclopedia because contributions would be publicly announced. The encyclopedia could knit together our profession in an im-

portant way, while fundamentally improving education and professional relationships in CS.

A nonprofit professional encyclopedia would be self-supporting through appropriate professionally relevant advertising carefully curated for high standards using existing advertising programs.

Vijay Kumar
Irrational Exuberance
and the 'FATE'
of Technology

[https://cacm.acm.org/blogs/blog-](https://cacm.acm.org/blogs/blog-cacm/230472-irrational-exuberance-and-the-fate-of-technology/fulltext)

[cacm/230472-irrational-exuberance-and-the-fate-of-technology/fulltext](https://cacm.acm.org/blogs/blog-cacm/230472-irrational-exuberance-and-the-fate-of-technology/fulltext)
August 20, 2018

I am sure many of us remember the Netscape IPO in 1995 and the fivefold growth in share value in four months. Expectations for technology and its impact were in the stratosphere. The Federal Reserve Board's then-chairman, Alan Greenspan, gave a speech at the American Enterprise Institute questioning "irrational exuberance" in the market and in technology.¹ I believe today we are seeing similar exuberance with technology.

Are revolutionary technologies for cancer screening—that rely on a fingerprick drawing one-thousandth the normal amount of blood—really feasible? Theranos had everyone believe such a revolutionary advancement was possible² not because of new techniques in analytical chemistry, but because it had developed novel software and new automation technologies! Can we really hope to replace eight million cars in Los An-

geles by boring tunnels³ for high-speed pods that will travel at 150 MPH at \$1 per ride? This is what Boring Company is selling the City of Los Angeles. Do recent advances in data science and machine learning really mean artificial general intelligence is around the corner? This is the pitch of so many startups today.

There have been advances in statistical machine learning, which have had remarkable impact in fields like computer vision and speech recognition when the underlying neural networks are trained by large-enough representative datasets. What “large enough” means, we don’t yet know. Neither do we know when we have a representative dataset. Yet there are many interesting cases where deep learning “works.” These success stories are oversold. In my own field—robotics—autonomy is a challenging problem, especially in tasks of manipulation and perception-action loops. Yet despite the claims being made, our best robots lack the dexterity of a three-year-old child.

Nowhere is irrational exuberance more evident than in self-driving cars. Not many people know the first demonstrations of an autonomous car were in the late 1980s at the Bundeswehr University Munich and at Carnegie Mellon University. Autonomous vehicles can have a tremendous social, economic, and environmental impact. This fact, and the technical challenges in realizing a bold vision, has attracted some of the top talent in science and engineering over the last 30 years. However, many of us don’t remember history, and many choose to ignore it since problems known to have not been solved for three decades are unlikely to attract investment.

According to recent predictions,⁴ fully autonomous cars will be available soon. Fully autonomous Audis and Teslas were promised several years ago by 2018. Uber even promised us flying cars powered by clean energy by 2023, even though the basic physics and chemistry underlying battery technology tells us otherwise.⁵

It is worrisome when engineers make these claims, and even more so when entrepreneurs use such claims to raise funding. However, the biggest concern should be about embedding software for autonomy in safety-critical systems. There is a difference between running

tests and logging data, and verification of software guaranteed not to have unwanted, unsafe behaviors. Can we claim vehicles are safe because the underlying software has been tested with over a billion miles of data? U.S. National Safety Council statistics suggest a billion miles of human driving, on average, results in 12.5 fatalities,⁶ and a billion-mile dataset cannot possibly be viewed as large enough or representative enough to train software to prevent fatalities.

The Uber-Waymo trial led to the release of documents truly shocking in this regard. They reveal a culture⁷ that appears to prioritize releasing the latest software over testing and verification, and one that encourages shortcuts. This may be acceptable for a buggy operating system for a phone that can be patched later, but should be unacceptable for software that drives a car.

This irrational exuberance may have its roots in the exponential growth in computing and storage technologies predicted by Gordon Moore five decades ago. The fact that just over a decade ago smartphones, cloud computing, and ride-sharing seemed like science fiction, and technologies like 3D printing and DNA sequencing were prohibitively expensive, has led to a culture of extrapolation fueled by exponential growth. Advances in creating programs that can play board games like chess and recent results with Alpha Go and Alpha Zero have been mind-boggling. Unfortunately, from this comes the extrapolation that it is only a question of time before we conquer general intelligence.

There is at least one argument that we are not making significant progress in understanding intelligence if we take into account the exponential growth in computing due to Moore’s Law. While computers have achieved superhuman performance in chess, the Elo rating of chess programs has merely increased linearly over the last three decades.⁸ If we were able to exploit the benefits of Moore’s Law, our chess-playing programs should be a billion times better than the programs from 30 years ago, instead of merely 30 times better. This suggests the exponential growth of technology may not even apply to algorithmic advances in artificial intelligence,⁹ let alone to advances in energy storage, biotechnology, automation, and manufacturing.

Irrational exuberance in technology has led to an even bigger problem: intellectual dishonesty, which every engineer and computer scientist must guard against. As professionals, it is our responsibility to call out intellectual dishonesty.

Questions of verification, safety, and trust must be central when we embody intelligence in physical systems. Questions of fairness, accountability, transparency, and ethics (FATE) should be addressed for data and information in society. It is great to see such efforts taking shape in industry¹⁰ and academia.¹¹

As teachers, we have an even bigger responsibility, as technology is no longer taught to just computer scientists or engineers; it is now a new liberal art. It is critical to address the true limitations of what technology can really bring about in the imminent future and the real dangers of extrapolation. Every university student who designs or creates anything must be sensitized to fundamental concerns of accountability and transparency and ethical responsibilities. We must address the FATE of technology, across all activities of design, synthesis and reduction of technologies to practice.

References

1. <https://www.federalreserve.gov/boarddocs/speeches/1996/19961205.htm>
2. <https://www.newyorker.com/magazine/2014/12/15/blood-simpler>
3. <https://www.cnn.com/2018/05/18/elon-musk-promises-1-rides-for-boring-companys-la-tunnels.html>
4. http://www.driverless-future.com/?page_id=384
5. <https://qz.com/1243334/the-magical-battery-uber-needs-for-its-flying-cars/>
6. <https://www.forbes.com/sites/jimgorzelany/2017/02/16/death-race-2017-where-to-find-the-most-dangerous-roads-in-america/#68a5354a1324>
7. <https://www.theverge.com/2018/3/20/17144090/uber-car-accident-arizona-safety-anthony-levandowski-waymo>
8. <https://www.eff.org/ai/metrics>
9. <https://freedom-to-tinker.com/2018/01/04/singularity-skepticism-2-why-self-improvement-isnt-enough/>
10. <https://www.microsoft.com/en-us/research/group/fate/>
11. <http://warrencenter.upenn.edu/>

Carl Hewitt is an emeritus professor of the Massachusetts Institute of Technology, USA. He is board chair of iRobust, a scientific society for the promotion of the field of Inconsistency Robustness, and board chair of Standard IoT, an international standards organization for the Internet of Things. **Vijay Kumar** is Nemirovsky Family Dean of Penn Engineering with appointments in the departments of Mechanical Engineering and Applied Mechanics, Computer and Information Science, and Electrical and Systems Engineering at the University of Pennsylvania, USA.